

1

**METHOD FOR DIVIDING COMPUTATIONS****CROSS-REFERENCE TO RELATED APPLICATIONS**

This application is a continuation of application Ser. No. 11/222,470, filed Sep. 8, 2005, now U.S. Pat. No. 7,243,121, which is a divisional of application Ser. No. 10/177,597, filed Jun. 20, 2002, now U.S. Pat. No. 7,103,628; the entirety of the foregoing related applications are incorporated by reference herein.

**BACKGROUND****I. Field of the Invention**

The present invention relates to the structure and operation of computing systems, and more particularly, to distributed computing systems and methods of operating such systems.

**II. Description of the Related Art**

Certain organizations have a need for high performance computing resources. For example, a financial institution may use such resources to perform risk management modeling of the valuations for particular instruments and portfolios at specified states of the world. As another example, a pharmaceutical manufacturer may use high performance computing resources to model the effects, efficacy and/or interactions of new drugs it is developing. As a further example, an oil exploration company may evaluate seismic information using high performance computing resources.

One conventional computing system includes a mainframe computer attached to an individual user terminal by a network connection. Using the terminal, a user may instruct the mainframe computer to execute a command. In this conventional system, almost all data storage and processing functionality resides on the mainframe computer, while relatively little memory or processing capability exists at the terminal. This terminal/mainframe architecture may not, however, allow computations requested by a user to be computed rapidly or automatically.

The open systems interconnection (OSI) model describes one conceptual network architecture represented by seven functional layers. In this model, the functions of a networking system in a data communications network are reflected as a set of seven layers, including a physical layer, data link layer, network layer, transport layer, session layer, presentation layer and application layer. One or more entities within each layer implement the functionality of the layer. Each entity provides facilities for use only by the layer above it, and interacts directly only with the layer below it. FIG. 1 depicts the seven functional layers of the OSI model.

The physical layer describes the physical characteristics of hardware components used to form a network. For example, the size of cable, the type of connector, and the method of termination are defined in the physical layer.

The data link layer describes the organization of the data to be transmitted over the particular mechanical/electrical/optical devices described in the physical layer. For example, the framing, addressing and check summing of Ethernet packets is defined in the data link layer.

The network layer describes how data is physically routed and exchanged along a path for delivery from one node of a

2

network to another. For example, the addressing and routing structure of the network is defined in this layer.

The transport layer describes means used to ensure that data is delivered from place to place in a sequential, error-free, and robust (i.e., no losses or duplications) condition. The complexity of the transport protocol is defined by the transport layer.

The session layer involves the organization of data generated by processes running on multiple nodes of a network in order to establish, use and terminate a connection between those nodes. For example, the session layer describes how security, name recognition and logging functions are to take place to allow a connection to be established, used and terminated.

The presentation layer describes the format the data presented to the application layer must possess. This layer translates data from the format it possesses at the sending/receiving station of the network node to the format it must embody to be used by the application layer.

The application layer describes the service made available to the user of the network node in order to perform a particular function the user wants to have performed. For example, the application layer implements electronic messaging (such as "e-mail") or remote file access.

In certain conventional high performance computing systems designed using the OSI model, the hardware used for computation-intensive processing may be dedicated to only one long-running program and, accordingly, may not be accessible by other long running programs. Moreover, it may be difficult to easily and dynamically reallocate the computation-intensive processing from one long running program to another. In the event processing resources are to be reallocated, a program currently running on a conventional high performance computer system typically must be terminated and re-run in its entirety at a later time.

**SUMMARY OF THE INVENTION**

In one aspect, the invention features a method including receiving, for computation by a node computing device of a distributed computing system, a parent job configured to produce one or more descendant jobs, wherein the node computing device is one of a plurality of node computing devices of the distributed computing system. Such a method also includes scheduling computation of the parent job on the node computing device, selectively rescheduling computation of a job other than the parent job from any one of the plurality of node computing devices to another of the node computing devices, and preventing rescheduling of the parent job unless each of the descendant jobs is completed or terminated.

**BRIEF DESCRIPTION OF THE DRAWINGS**

Features and other aspects of the invention are explained in the following description taken in conjunction with the accompanying drawings, wherein:

FIG. 1 depicts the seven functional layers of the open systems interconnection (OSI) model;

FIG. 2 illustrates a system 10 including a compute backbone 300 according to one embodiment of the present invention;